

PARALLEL DATA INTERFACE

BACKGROUND OF THE INVENTION

The present invention relates to an interface or other apparatus receiving parallel transmitted data streams.

There are generally two well recognised ways in which data is transmitted. In serial data transmission the data is transmitted sequentially via a single transmission channel. In parallel transmission a plurality of associated channels are provided and data is transmitted simultaneously via the plurality of channels.

In any data transmission system, the data is generally transmitted in a fixed relationship to a clock signal. That is, a clock signal defines fixed time slots and one data bit is transmitted in each time slot. Upon reception of the transmitted signal the relationship of the received signal to the data time slots must be established to enable recovery of the transmitted data. Because of variations introduced by the transmission medium it is not possible simply to run a clock having an appropriate frequency at the receiver without ensuring that it is properly synchronised with the incoming data.

In serial transmission systems a suitably synchronised clock at the receive apparatus may be generated from the received data itself, or the data sequence may be used to synchronise a locally generated clock to enable data recovery. Using such arrangements, high data transmission rates have been achieved using serial data transmission technique.

Parallel data transmission presents other problems in terms of data recovery. In particular the transmission characteristics of each of the plurality of parallel channels are not always identical. Some variation may be introduced by the physical construction (e.g. cable lengths) of the transmission paths and these can be minimised by appropriate design. Other factors include interference in the path and it happens that such environmental factors affect some channels differently to others. One effect of these different characteristics in the various channels is that the transmission time from transmission to reception may not be identical for all channels. Thus, at the receive apparatus there may be some departure from proper synchronisation between the channels and this is known as Askew@.

Typically, one channel in a parallel system may be used to transmit a clock signal which can be used for the data recovery at the receiver, and the skew also affects the timing relationship between the clock channel and the data channels.

It is possible to avoid errors caused by skew in a parallel transmission system between the data channels and the clock channel by taking steps such as limiting the transmission distance and the data rate in each channel. This has the effect that the magnitude of the skew introduced is small compared to the data clock intervals, so that it does not interfere with the data recovery.

However, as bandwidth requirements in data transmission systems increase there is demand for the ability to transmit parallel data at data rates in each channel approaching those previously used for serial transmission. At such data rates the problems caused by

skew in the parallel transmission channel have a significant effect in the ability to recover received data.

One approach would be to regenerate a separate data recovery clock for each of the parallel channels. This is however impractical for a large number of parallel data channels, and also does not deal with the lack of synchronisation between the data channels.

SUMMARY OF THE INVENTION

The present invention provides apparatus for receiving parallel transmitted data in a plurality of channels comprising means to generate a clock signal on the basis of the received data and means associated with each of said channels to synchronise data received on the associated channel with the generated clock.

In this arrangement a single clock signal is generated which is used for all the data channels. This means that the apparatus is easily scaleable to receive data from large numbers of parallel channels.

In synchronising all the data channels with a single clock the apparatus also removes the skew between the data channels. Thus the apparatus can simply present as-received but re-aligned data signals for subsequent processing. Alternatively the apparatus can perform the data recovery at the same time as re-aligning the channels.

The clock signal may be generated on the basis of a single received channel. That channel may be a channel designated for the transmission of a clock signal from the transmitter. Alternatively, that channel may be

one of the data channels in which it is expected that there will be a significant number of data transitions.

It may also be possible to generate the clock signal on the basis of a plurality of the parallel channels.

The synchronising of each data channel with the clock is preferably done by applying a variable delay to each of the data channels. Also, the generated clock signal is preferably delayed by half the maximum delay available to each data channel so that the data channels can be effectively advanced or retarded in relation to the clock.

BRIEF DESCRIPTION OF THE DRAWINGS

The problems overcome by the invention together with other features and advantages will be more fully explained in the following description of a preferred embodiment, given by way of non-limiting example, and with reference to the accompanying drawings, in which:

Figure 1 shows ideal clock and data signals;

Figure 2 shows clock and data signals with skew;

Figure 3 shows an outline of high-speed parallel interface;

Figure 4 illustrates example phase detector with idealised signal waveforms;

Figure 5 shows a phase detector characteristic;

Figure 6 shows a phase detector characteristic with data delay adjustment range for ideally aligned data;

Figure 7 shows a phase detector characteristic with data delay adjustment range for misaligned data;

DRAFTS & DRAWINGS

Figure 8 shows a phase detector characteristic with high skew and large T_d ;

Figure 9 shows a phase detector characteristic with high skew and large T_d with delay Awrap around@;

Figure 10 illustrates a variable data delay line based on interpolator; and

Figure 11 illustrates an extended data phase interpolator delay line for improved linearity/range.

DETAILED DESCRIPTION

Figure 1 illustrates data signal timing in a typical data transmission system. In particular Figure 1 shows a clock signal 10, known as a half-rate clock, and data slots are defined between clock transitions. This is shown by the representative data stream 12 with sequential data slots 14. In the preferred embodiment it will be assumed that a half-rate clock is transmitted in one of the parallel channels. For data recovery it is usual to re-generate a full-rate clock having a frequency twice that of the half-rate clock which therefore has transitions in the centres of each data slot 14 as well as at the boundaries.

Figure 2 is a diagram similar to Figure 1 but illustrating the effect of skew in the transmission channel. As compared to half-rate clock 10 it can be seen that the boundaries between the data slots in data stream 22 can drift from synchronisation with the clock transitions as a result of variations in the transmission times in the various channels.

More precisely, skew is specified by a single time value representing the maximum alignment error

between any two signals in the parallel transmitted signals. This is defined as T_s and, at worst therefore, any data bit could be shifted early or late with respect to the clock by up to T_s . This receiver needs to be designed to handle such misalignment.

An outline of the preferred parallel interface receiver system is shown in Figure 3. This comprises a Clock Recovery circuit 30 and a set of Data De-skew circuits 40, one for each bit in the parallel bus. The basic principle of the system is to generate a recovered clock from the Clock input and to distribute this to each of the de-skewing circuits where each of the incoming data signals is shifted into alignment with the clock using a variable delay line.

The operation of the De-skew circuits 40 will be described in more detail below, but it may be noted that each such circuit comprises a variable delay 42 which is arranged to apply a variable delay between 0 and T_d to the received data. The delay 42 is controlled by a delay line control means 44 which operates on the basis of a comparison between the delayed data and the clock signal effect by phase detector 46.

A delay line 32 is also used in the clock recovery circuit 30, where it is set to give a delay exactly in the middle of its range: ie the delay line in the clock recovery block is set to $2T_d$. This allows the data to be shifted with respect to the clock by $\pm 2T_d$ in the data de-skewing blocks 40.

The clock recovery system shown is based on a phase interpolation technique wherein an output clock phase is generated from a pair of quadrature reference clocks 35 by summing these with different weightings in a

phase interpolator 34. In Figure 3, the reference clocks (and hence the aligned data clock) will nominally be at the full data rate. However, it is possible to adapt the system to operate on a half-rate clock. Control of the phase interpolator 34 is performed using a phase detector 38 to compare the alignment of the recovered clock 50 with the delayed half-rate clock. This then produces control signals which are used to adjust the phase interpolator weightings. The phase interpolator control 36 is generally carried out using digital techniques, although the analogue method described in patent application 0004298.6 may also be used.

The recovered clock 50 is distributed to each of the data channels. In practice, care needs to be taken to ensure that this clock distribution does not itself exhibit skew. The data de-skewing circuits 40 then use phase detectors 46 which may be identical to that in the clock recovery block 30 to control the variable delay lines 42 so as to shift the data into alignment with the recovered clock 50.

The delay lines allow the data to be shifted in position with respect to the clock by $\square 2T_d$, therefore in order to ensure that the skew can be cancelled out at each input it must be ensured that $2T_d > T_s$.

The precise implementation of the phase detector 38,46 is not a part of this invention. However, in general this will simply provide an indication to either increase the delay (via the "Up" control signal) or decrease the delay (via the "Down" control signal) if the data is early or late respectively. A simple example of a possible phase detector circuit 46 is shown in Figure 4A. This circuit simply samples the received data on the

FIGURE 3 - FIGURE 3

positive and negative edges of the clock 50 by way of latches 402, 403. Exclusive-OR function 404 detects changes in the data value: if the change occurs between a positive clock edge and the ensuing negative edge it is considered early and an "Up" pulse is generated by latch 405, whilst if the change occurs between a negative clock edge and the ensuing positive edge it is considered late and a "Down" pulse is generated by latch 406. In this way, the data edges are brought into alignment with the negative clock edges, and therefore the positive clock edge of the full-rate clock is centred in the data eye to optimally sample the data bit values. This timing is illustrated in Figure 4B.

This phase detector behaviour can be described by the characteristic shown in Figure 5. Note that this characteristic exhibits a periodicity bounded by $\Delta 2UI$, where UI is a "unit interval" which is equivalent to the period of a single data bit. This is a necessary characteristic of a data phase detector.

In the de-skewing circuits 40, the phase detector 46 is used to control the data input delay line to adjust its phase with respect to the aligned data clock 50. Figure 6 shows the adjustment range ($\Delta 2T_d$) of the data signal for an ideally aligned input superimposed onto the phase detector characteristic. Figure 7 shows a similar diagram for misaligned data: in this case, the data is late and the phase detector will indicate that the delay needs to be reduced. This diagram illustrates the earlier stated condition; that in order to re-centre the data, $2T_d > T_s$.

Figure 8 shows a similar diagram to Figure 7, but with a higher value of skew and a correspondingly

increased data delay adjustment range. Under these conditions, it is possible to adjust the phase of the data such that it overlaps into the adjacent bit period. If the system were to get into this state, the phase detector 46 would indicate the wrong direction to centre the data (e.g in Figure 8, the phase detector would try to increase the delay rather than reduce it) and would potentially lock up at the end stop of the delay line range. It can be seen that the condition for this to occur is that $T_s + 2T_d > 2UI$.

The range for T_d to meet these requirements is therefore as follows:-

$$T_s < 2T_d < (2UI - T_s)$$

These constraints could prove a serious limit to the practicality of this system in reality, since T_d will be subject to variation due to manufacturing tolerances, whilst any increase in T_s results in a decrease in the tolerable range of T_d for both its minimum and maximum values. For instance, if $T_s = 3UI$, T_d has zero margin for error.

In order to alleviate these constraints, it is desirable to avert the potential lock-up condition. In fact it is possible to do this by allowing the delay line control to wrap around from its maximum value to its minimum value and vice versa. If this is implemented, no potential lock-up will occur unless the skew and data delay are sufficient for it to lock onto the centre of the adjacent data bit as shown in Figure 9. This will only occur if $T_s + 2T_d > 2UI$. Thus our restrictions for T_d are now as follows:

$$T_s < 2T_d < (UI - T_s)$$

which gives considerably more margin than the previous case.

Note that the requirement to allow wrap around of the data delay lines will probably mandate a digital solution to control these.

Although there are various standard ways to implement the variable delay line, one preferred implementation is shown in Figure 10 and makes use of a fixed delay element 102 in conjunction with a variable interpolator 104. Phase interpolator 104 mixes the non-delayed signal D0 in variable proportions with maximally delayed signal D1 to output a variable delay signal. This may be implemented as illustrated by a pair of transistor pairs 106, 107 to which differential representations of D0 and D1 are applied and mixed in variable proportions according to the values of current sources I0, I1. In this scheme, the bias currents I0 and I1 are varied in opposition so that the total current is constant.

The design in Figure 10 provides good performance providing that T_d is relatively small compared with the data bit period. For higher values of T_d , the circuit of Figure 11 may be used, which provides a number of delay stages 112 rather than a single slow stage (which will tend to attenuate the high speed data signal components). These could then be used in conjunction with a multi-stage interpolator akin to that shown in Figure 10. The delay line could be further extended with a larger number of stages if required. This would tend to both improve linearity of the data phase interpolator and allow a larger delay variation.

CROSS-REFERENCED